Language, Ontology, and the Semantic Web

Reclaiming the Original Vision for a More Intelligent Web

eswc20

The Original Semantic Web Vision Was a Unifying Logic Bus, Not Just a Data Layer

- In 2000, Tim Berners-Lee proposed an ambitious vision for the Semantic Web as an "interchange bus for on-line data."
- The core of this vision was a powerful Semantic Web Logic Language (SWeLL) that would extend RDF.
- **Key Goals of SWeLL:**
  - Include negation and explicit quantification.
  - Represent first-order and even higher-order logic.
  - Connect simple, reliable systems with complex, heuristic ones.
- This 'unifying language for classical logic' was designed to be the central bus for all data and rules, enabling true interoperability and reasoning.

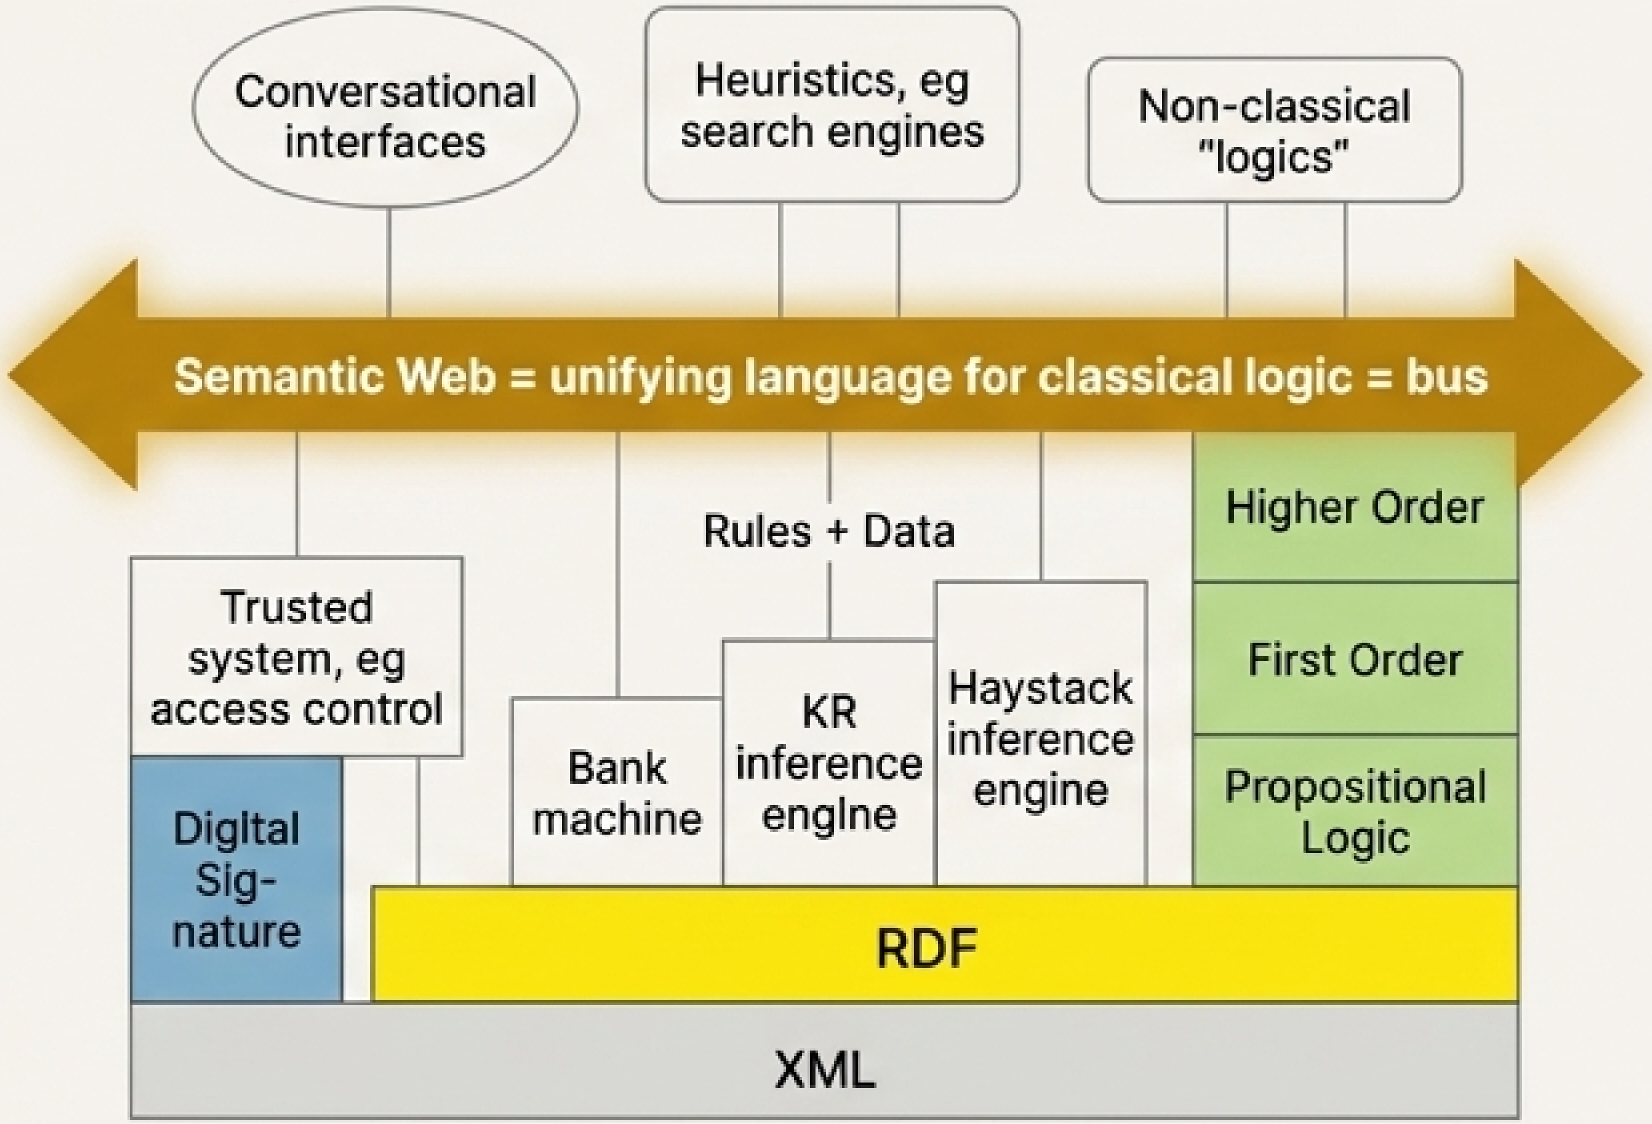

The 2005 Reality Delivered a Stack, Not a Bus, Sacrificing Expressiveness for Decidability

- The system delivered in the 2005 W3C Final Report diverged significantly from the original proposal.
- The unifying logic bus (SWeLL) was replaced by a ‘layer cake’ of standards like RDF-S and OWL.
- This new stack prioritized decidability (guaranteed to terminate) over logical expressiveness.
- **The Critical Omission**: The powerful, unifying logic layer capable of handling first-order logic and beyond was gone. The system could no longer serve as a universal interchange for diverse logical systems.

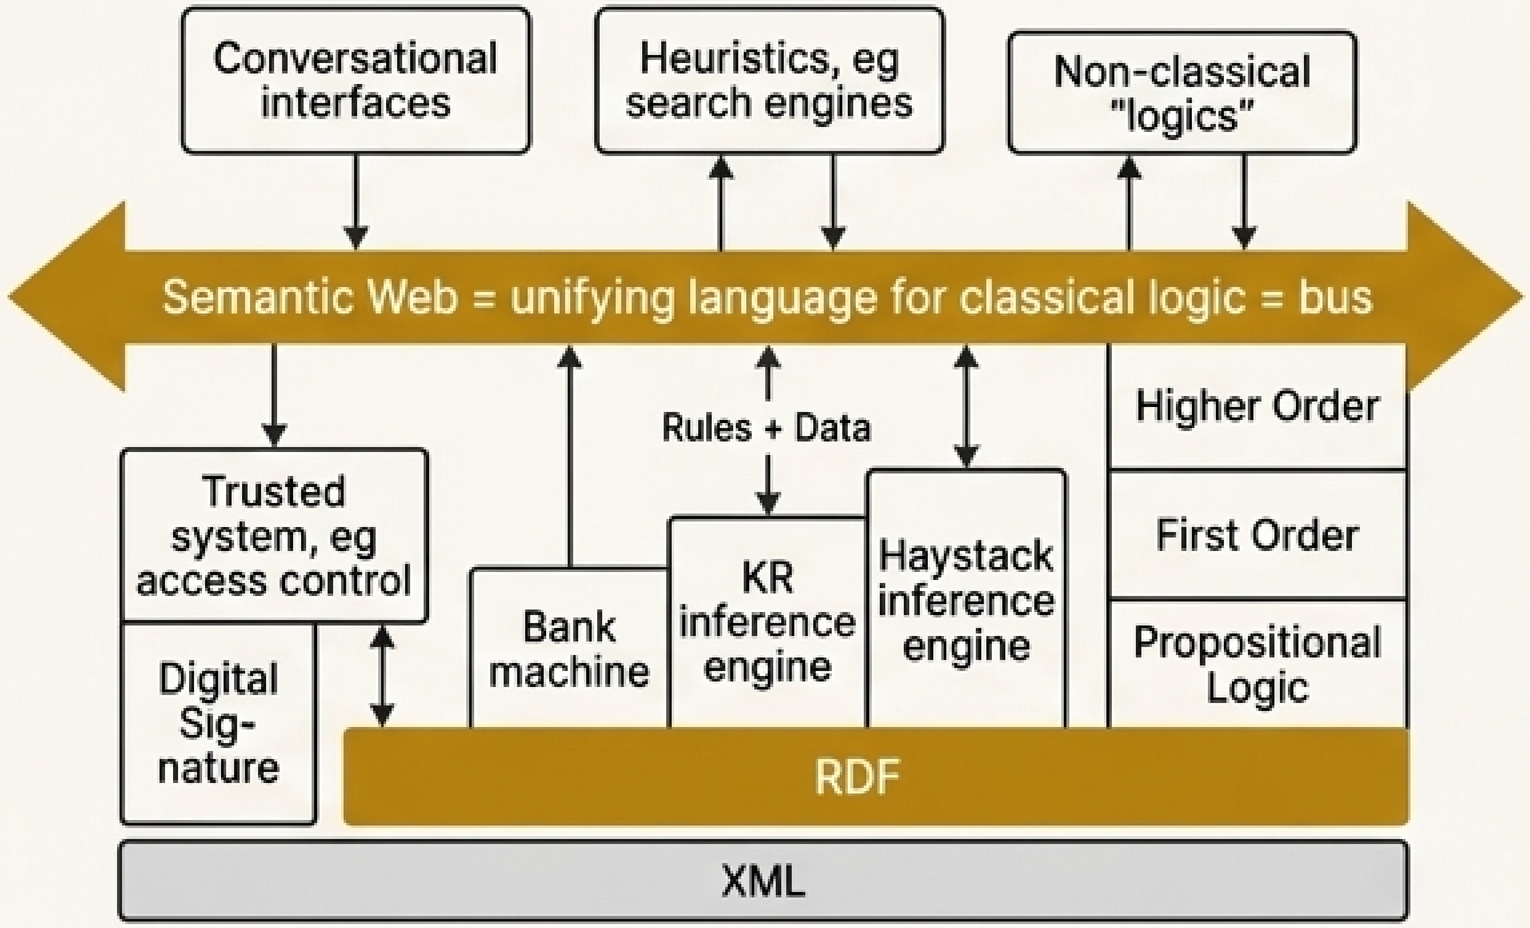

Winning Proposal (2000)

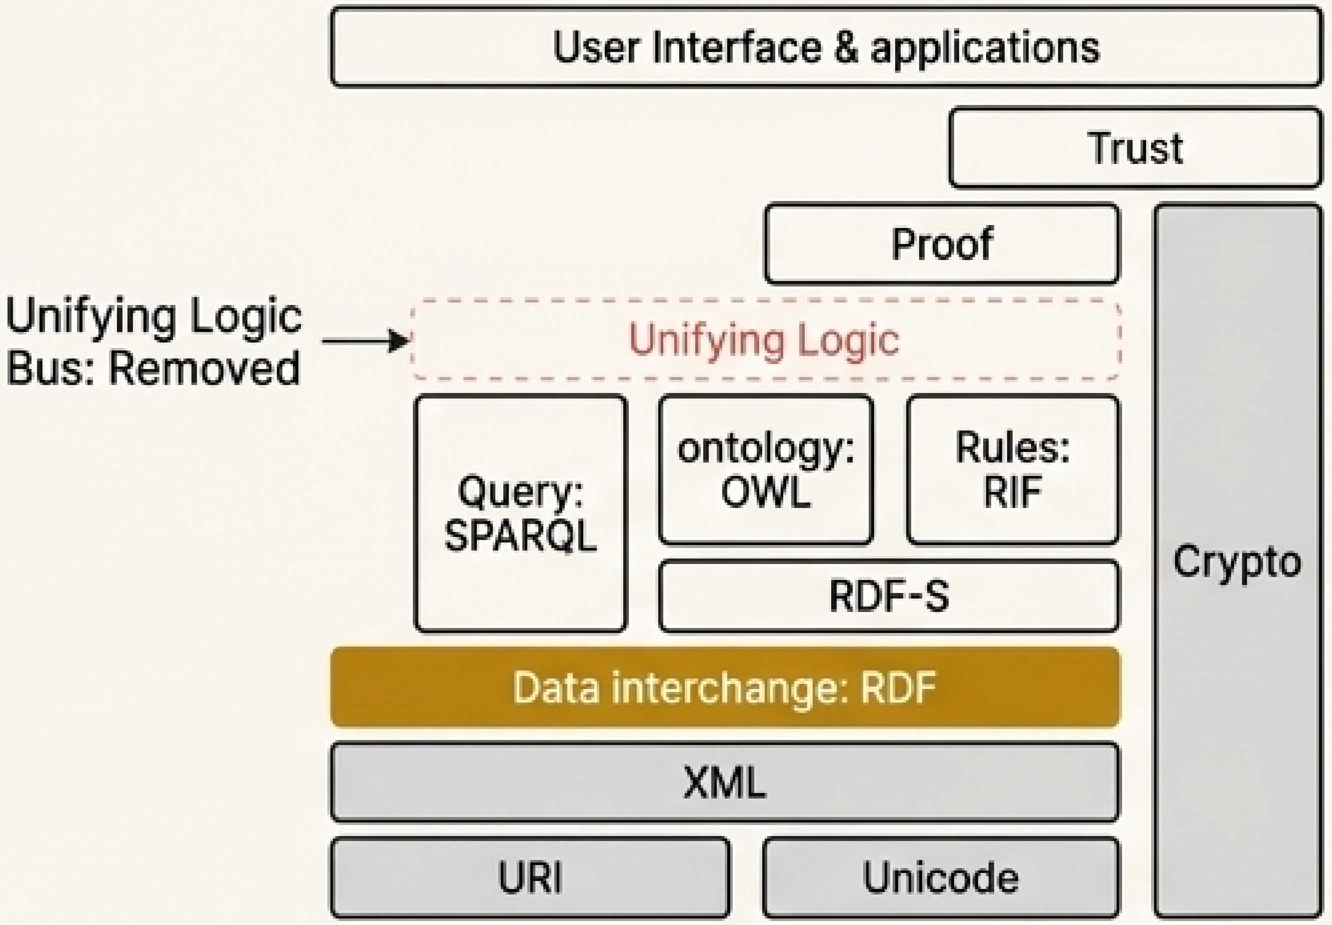

Final Report (2005)

The 'Decidability' Trade-off Was a Fallacy That Crippled the Semantic Web's Potential

- The Semantic Web logics (like OWL) were intentionally restricted to be decidable. But this was a misguided goal.
- **Key Argument:** Decidability is a property of a *problem*, not the *notation*. Powerful theorem provers (like TPTP systems) use syntactic checks to select efficient methods and are just as fast for the same class of problems.

DECIDABILITY

"Restricting expressive power cannot improve performance. It just makes certain problems impossible to state."

"Users always ask for more expressive power. They never ask for decidability." – John F. Sowa

The result was a system that couldn't represent the complex knowledge needed for true AI applications, while offering no real performance benefit for the problems it *could* solve.

Common Logic Fulfills the Original Vision as a Standardized Unifying Logic

- The original proposal for the Semantic Web Logic Language (SWeLL) evolved into the ISO/IEC standard 24707: **Common Logic (CL)**.
- CL is a highly expressive logic designed for information exchange and interoperability between disparate systems.
- It supports the full range of first-order logic, with extensions for metalanguage and higher-order reasoning, restoring the power of the original vision.
- **Introducing CLIP (CL Interface to Predicate calculus)**: A human-readable dialect of CL designed to be:
  - As readable as Turtle for RDF/OWL subsets.
  - A clear linearization for various graph logics (KGs, UML, etc.).
  - A bridge between formal logic and natural language.

"Common Logic is SWeLL."
– John F. Sowa

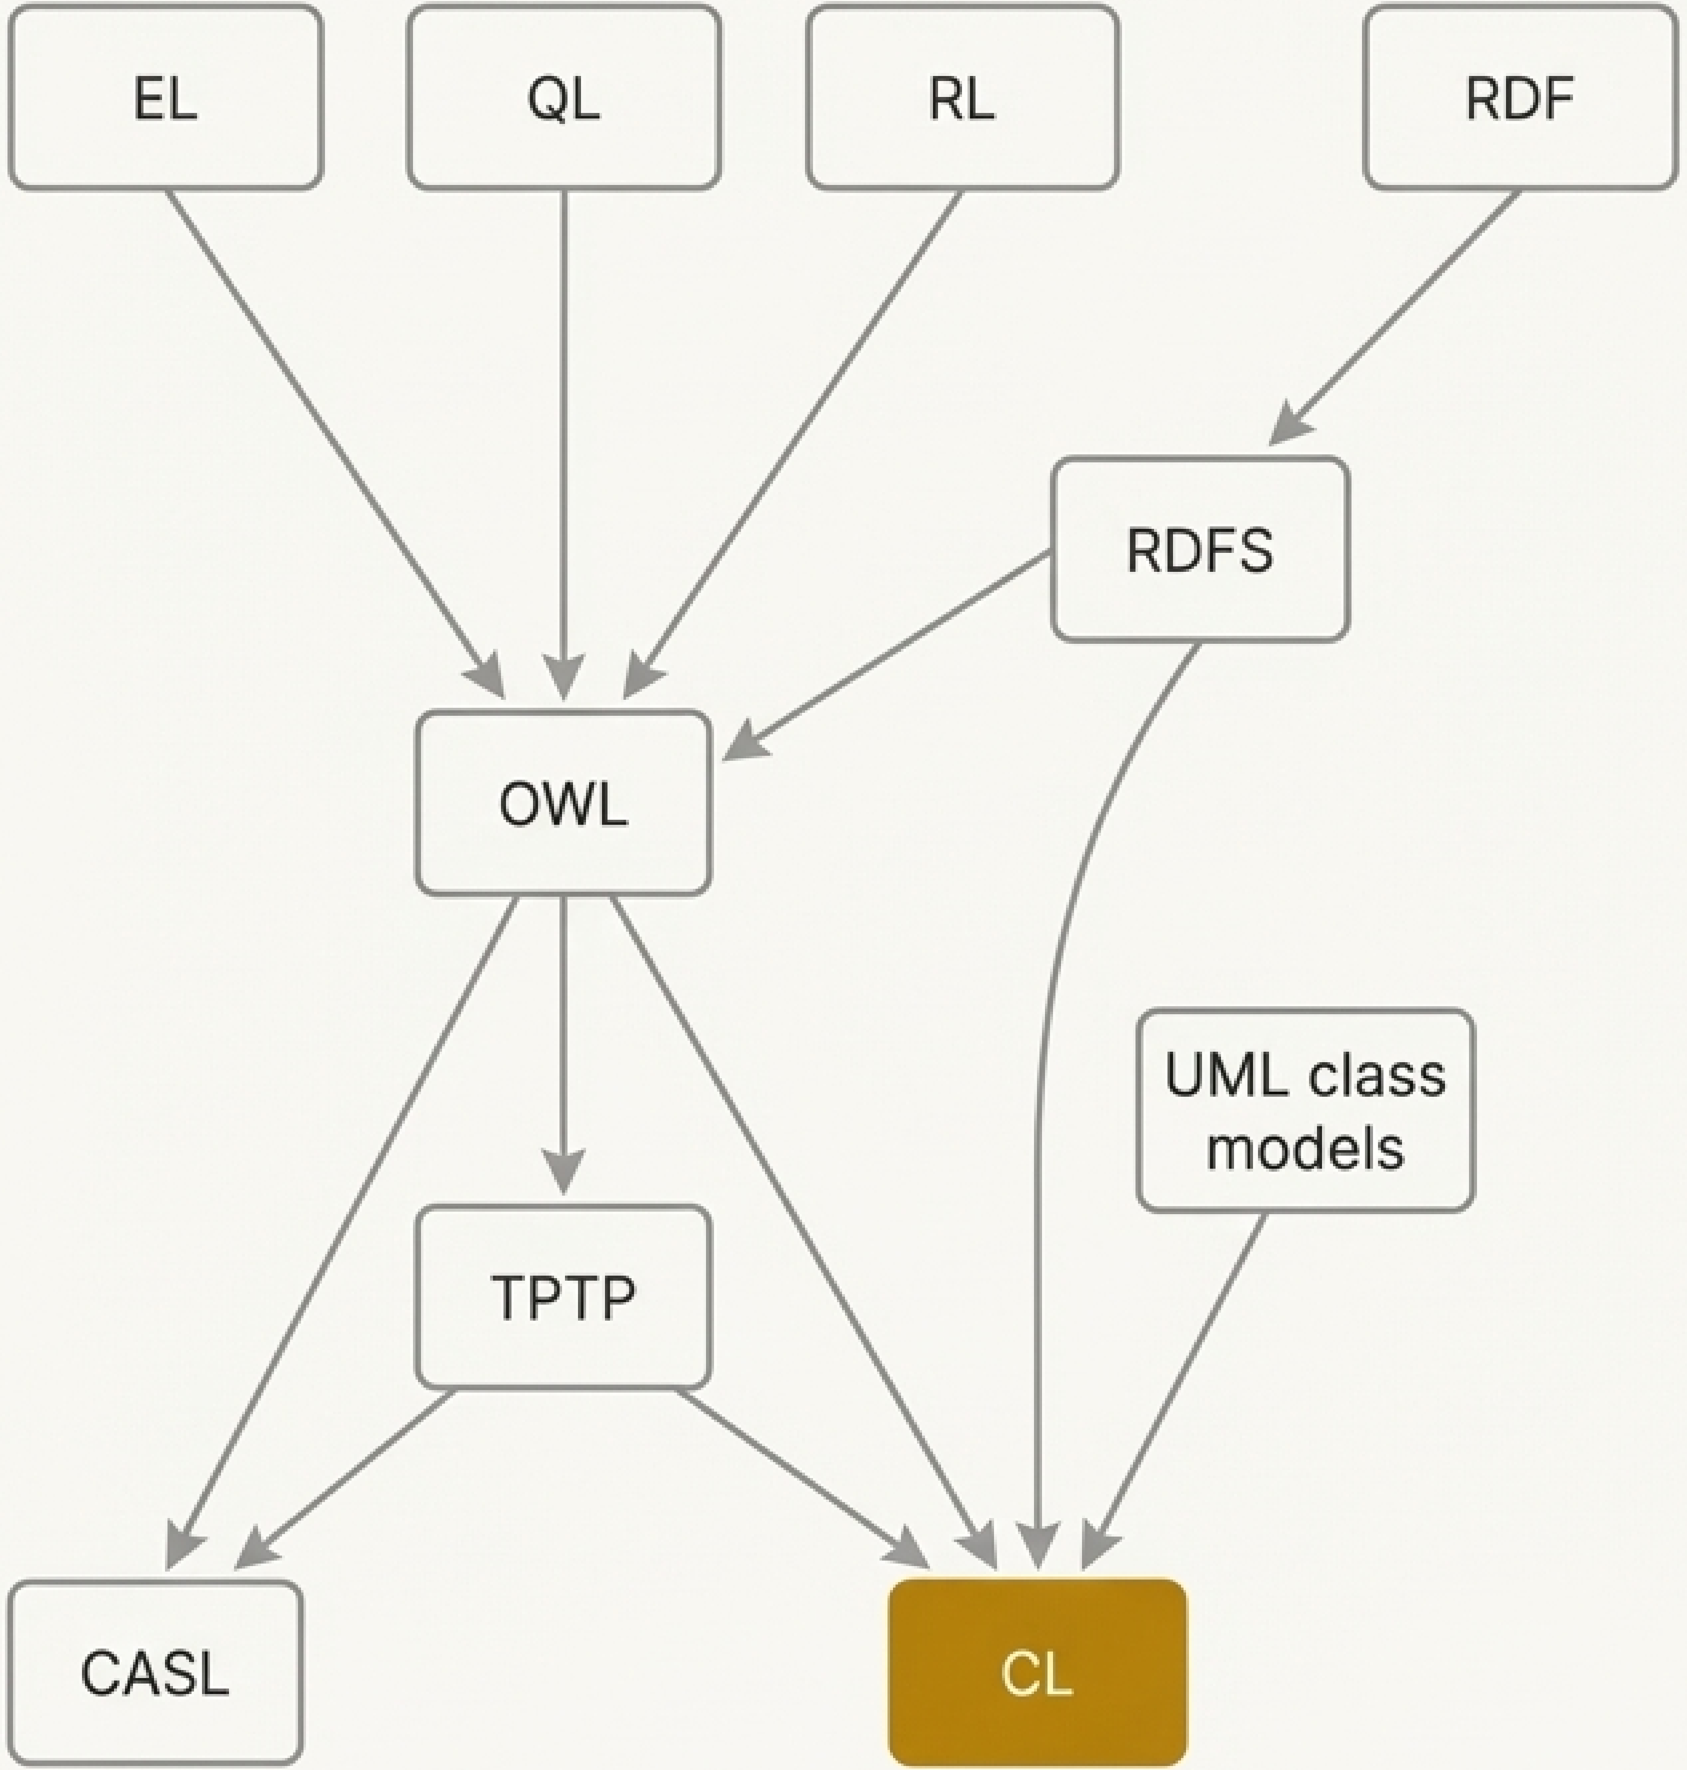

Common Logic Provides a Hub for True Interoperability Across Heterogeneous Systems

- The Distributed Ontology, Modeling, and Specification Language (DOL) standard demonstrates how to integrate diverse systems by relating their underlying logics.
- Common Logic is one of the most general logics supported by DOL, capable of serving as the interchange hub.
- By mapping languages like OWL, RDF, and even UML class models to CL, we can bridge the gaps that currently 'kill' productivity.

"Any one of those tools, by itself, is a tremendous aid to productivity, but any two of them together will kill you."
— Terry Rankin, 1980. CL is the solution.

CLIP: Expressing Full First-Order Logic with the Clarity of Natural Language

A simple sentence, “A cat is on a mat,” reveals the evolution and structure of logical notations.

Peirce (1897, Existential Graph)

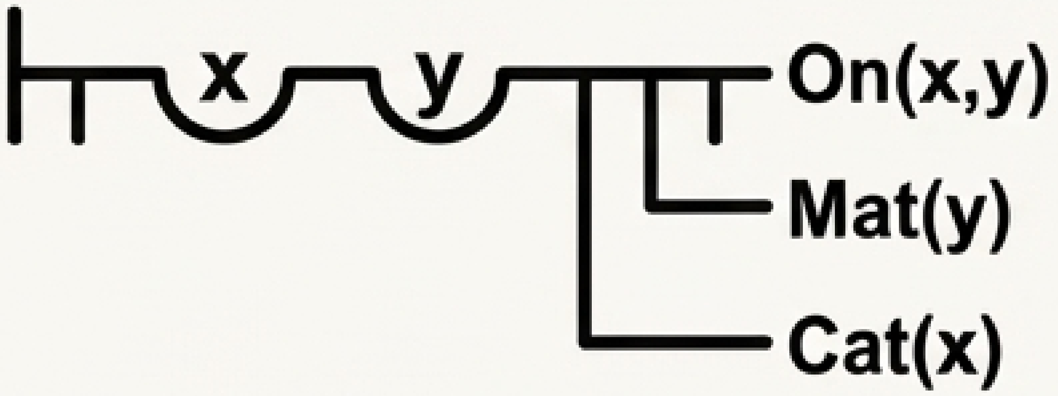

Peano (1895, Predicate Calculus)

$$\exists x\, \exists y\, Cat(x) \wedge On(x, y) \wedge Mat(y)$$

Common Logic (CLIP Dialect)

```
(Exists (x y) (Cat x)
  (On x y) (Mat y))
```

CLIP achieves the expressiveness of predicate calculus while maintaining a structure that is both human-readable and directly mappable to graph-based representations like Peirce’s Existential Graphs (EGs), which are considered more ‘iconic’ and cognitively natural.

Beyond Data: A Holistic Framework for Language, Logic, and Cognition

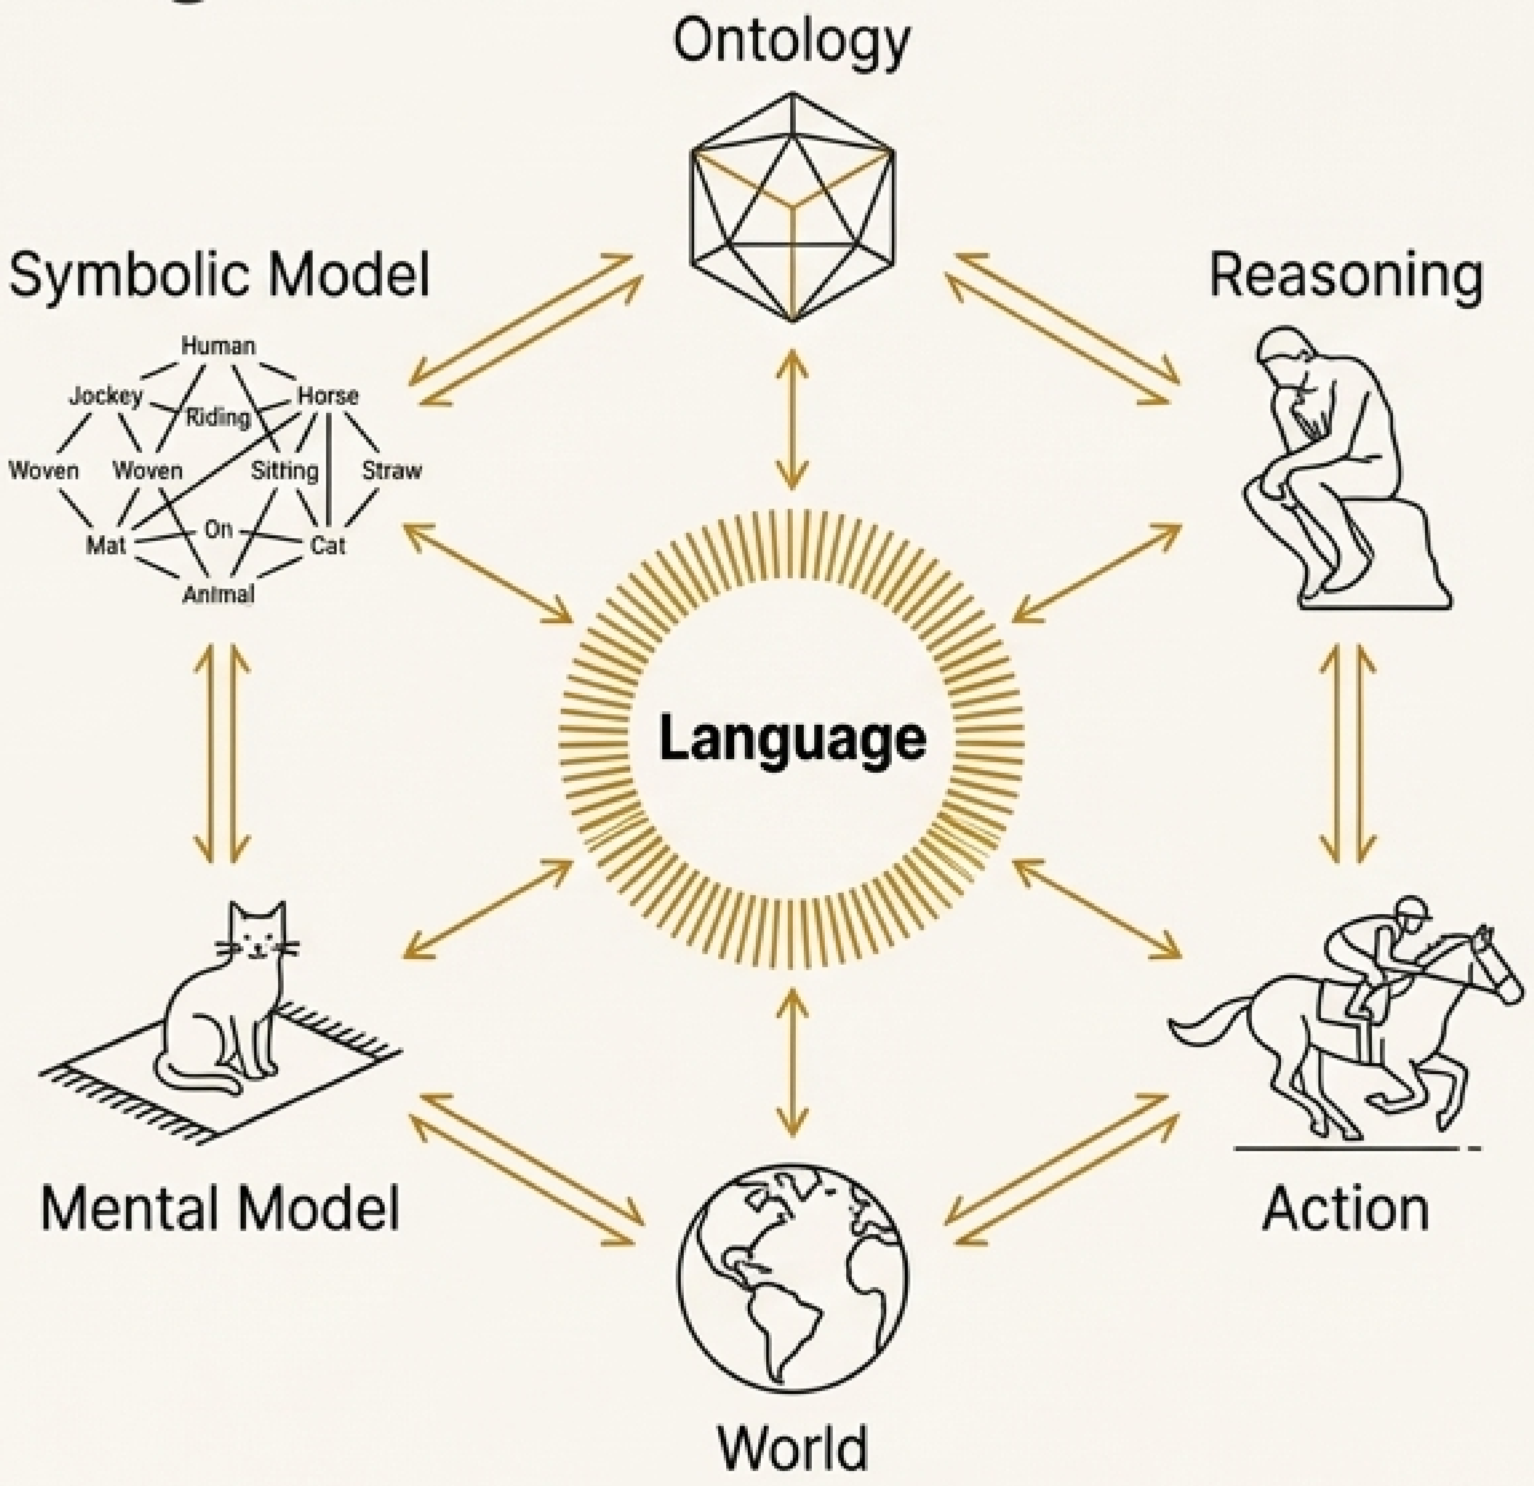

- To achieve true AI, we need to connect logic to a broader cognitive architecture. Sowa's Hexagon of Knowledge provides this map.
- **The Six Aspects of Knowledge:**
  1. **World**: What we encounter.
  2. **Mental Model**: What we experience or imagine.
  3. **Symbolic Model**: Words related to other words (e.g., KGs).
  4. **Ontology**: A catalog of words and what they refer to.
  5. **Reasoning**: How we think.
  6. **Action**: What our thinking leads us to do.
- **Central Insight**: Natural Language is the central element that can represent and relate all six corners. Any artificial language is a simplification of something sayable in natural language.

Current AI Excels at Perception; True Cognition Requires Causal Reasoning

- Most current Machine Learning methods excel at tasks that take humans less than a second—fast pattern recognition and classification (S-R theories).
- However, true cognitive learning, as studied in neuroscience, involves a progression of deeper understanding.
- **fMRI studies show learning progresses in stages**:
    1. **Visual Perception**: Recognizing shapes and parts (Occipital lobes).
    2. **Thinking about Structure:** Understanding mechanical and spatial relationships (Parietal lobes).
    3. **Thinking about Causality:** Generating hypotheses about how a system works and how one would interact with it (Frontal lobes, connections across the brain).
- This demonstrates a clear neurological basis for moving beyond pattern-matching to structural and causal models.

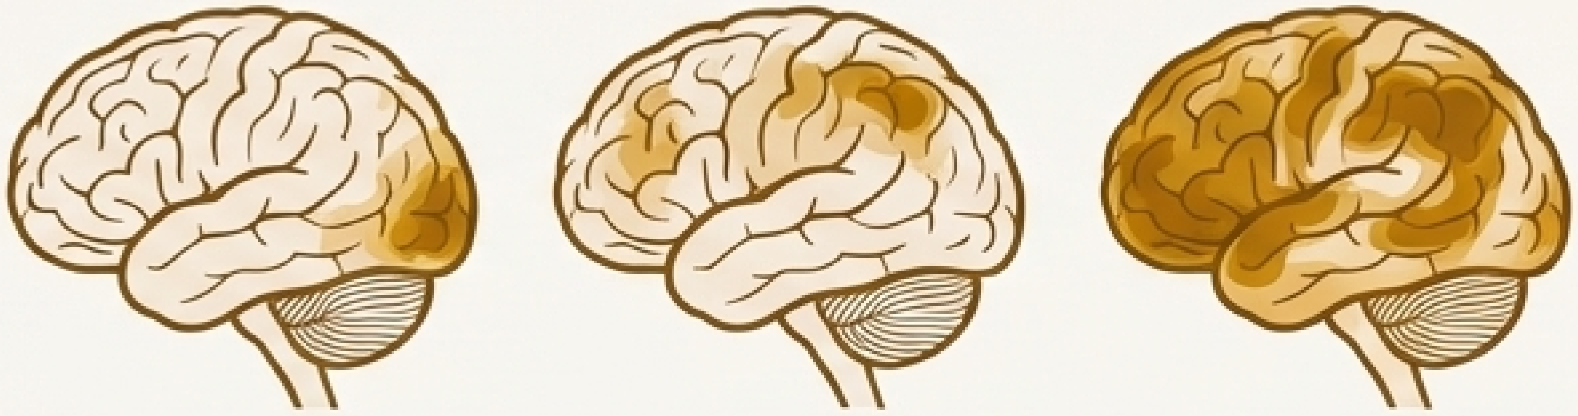

1. Visual perception

2. Thinking about structure

3. Thinking about causality

The Path Forward is a Neuro-Symbolic Hybrid Combining Two Modes of Reasoning

The human brain combines two complementary systems for reasoning about the world:

1. **Symbolic Models (Cerebrum/Frontal Lobes)*: Necessary for language, logic, and precise inference. Symbols provide abstraction and structure but have an indirect mapping to the world.

2. **Mental Models (Cerebellum)*': Essential for perception, action, and simulation. These models are grounded in sensory-motor experience and provide a direct, intuitive link to the world.

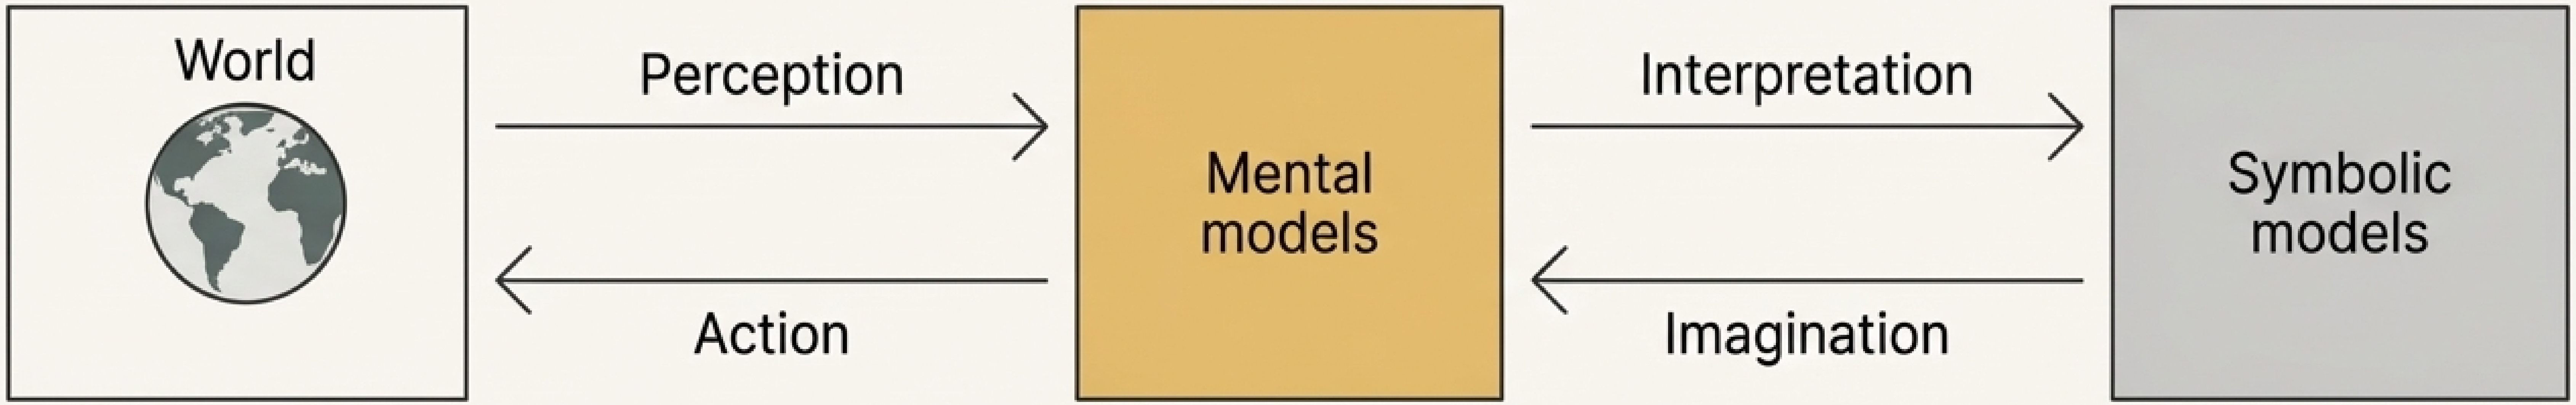

A hybrid AI architecture that integrates neural networks (for perception and simulating mental models) with symbolic systems (for language and logic) can combine the advantages of both, leading to more robust and human-like intelligence.

The Goal is Not Just a More Powerful System, but a More Intelligent One

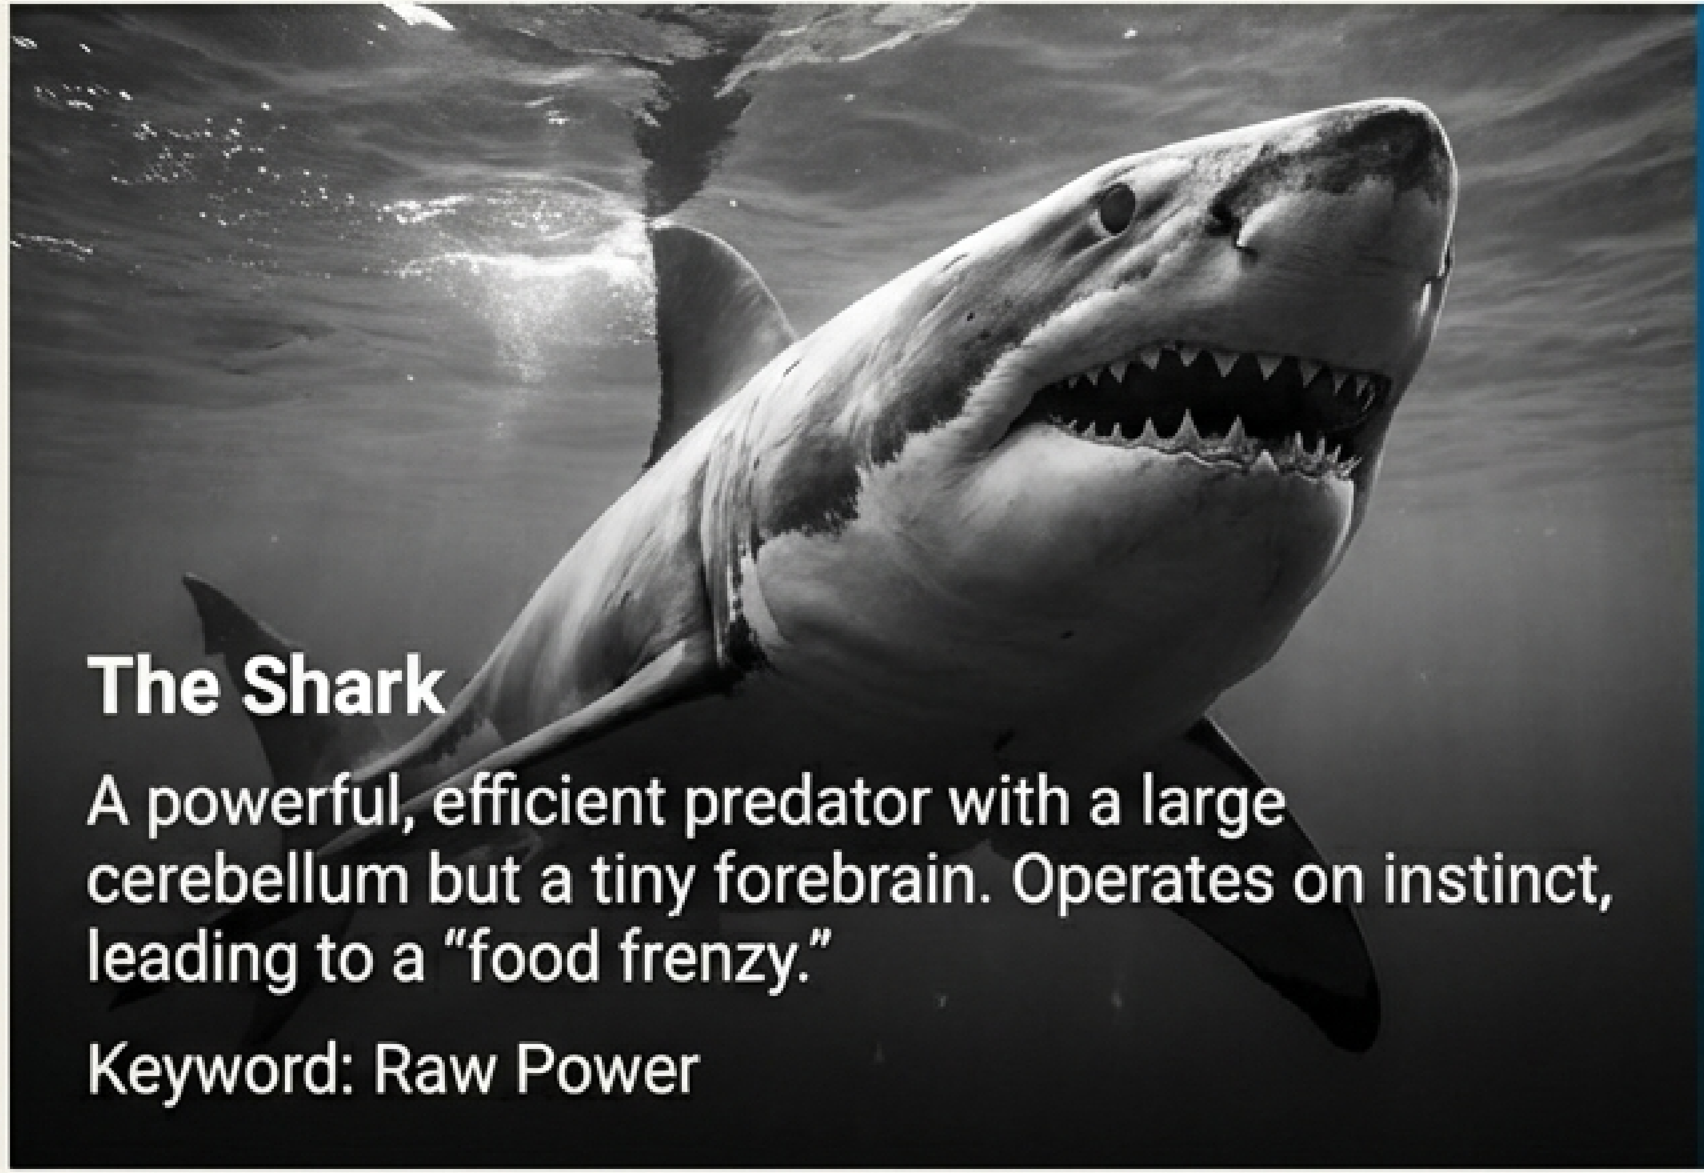

For AI, the goal should not be to build bigger sharks. The goal should be to build dolphins: systems that combine perceptual prowess with symbolic reasoning to achieve communication, coordination, and truly intelligent behavior.

A Practical Tool for Building Better Ontologies: Formal Concept Analysis (FCA)

- FCA is a semi-automated theory and toolset for designing, aligning, and merging ontologies.
- **How it Works**: FCA algorithms compute a minimal concept lattice from a set of types and their defining attributes, revealing all inheritance paths and ensuring consistency.
- **Key Applications for the Semantic Web**:
  - **Ontology Development**: Automatically generating lattices from lexical resources.
  - **Ontology Merging:** Detecting conflicts and similarities to integrate independently developed ontologies.
  - **Consistency Checking**: Verifying that ontologies specified in OWL and other notations are internally consistent.

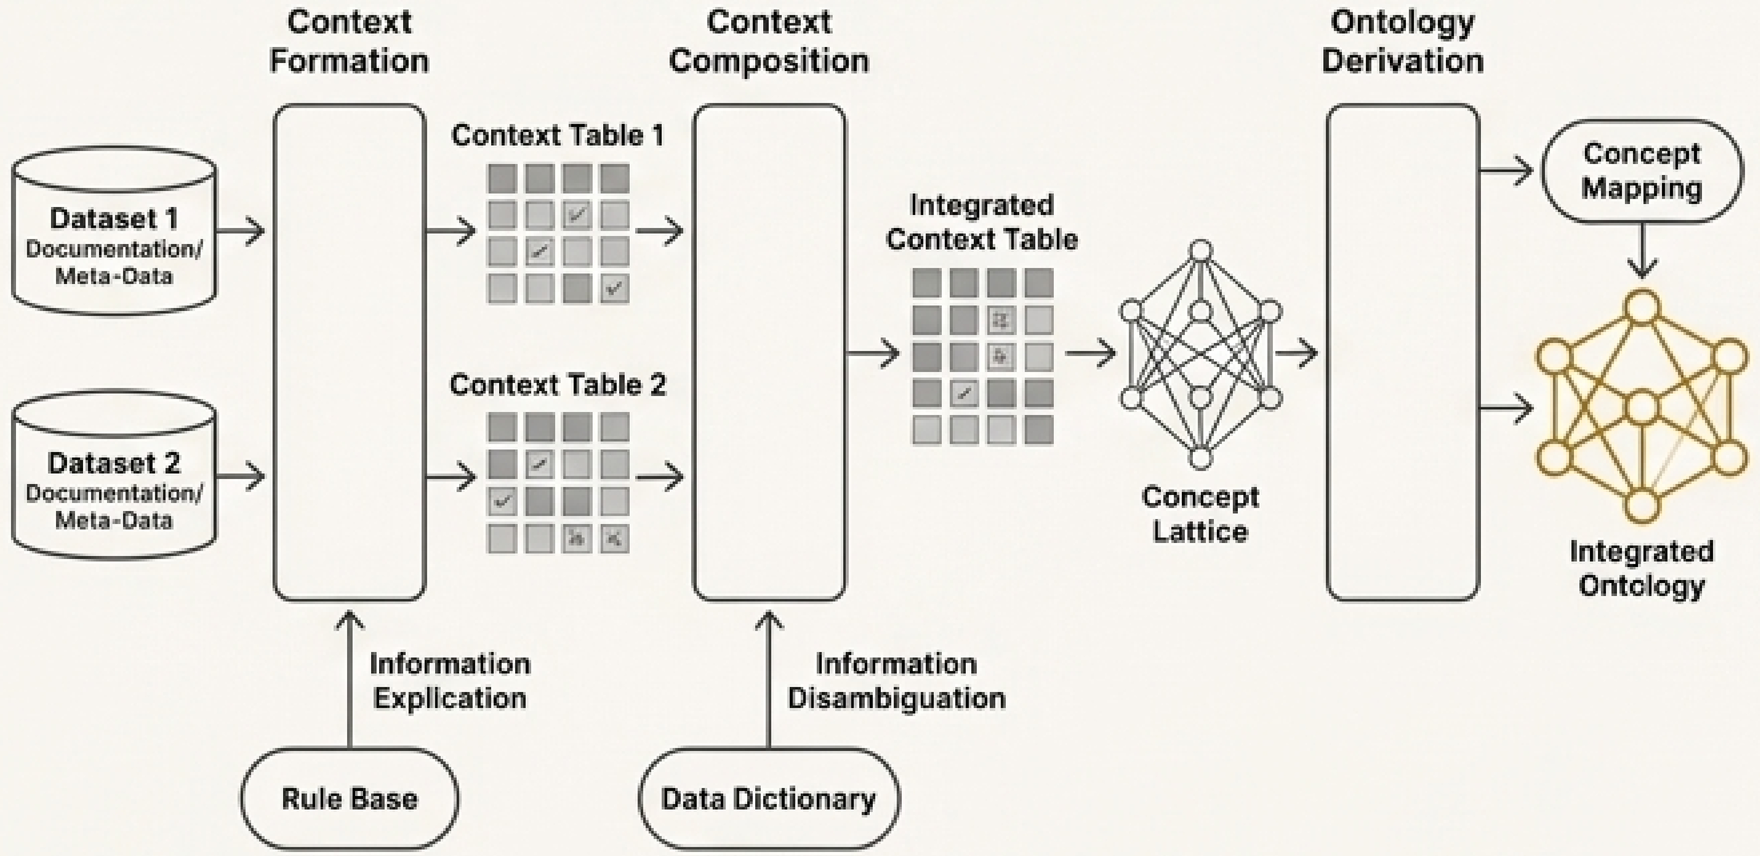

NotebookLM